Visualizing Cyber Security Risks with Bow-Tie Diagrams

Karin Bernsmed, Christian Frøystad, Per Håkon Meland, Dag Atle Nesheim and Ørnulf Jan Rødseth

Citation:
Bernsmed K., Frøystad C., Meland P.H., Nesheim D.A., Rødseth Ø.J. (2018) Visualizing Cyber Security Risks with Bow-Tie Diagrams. In: Liu P., Mauw S., Stolen K. (eds) Graphical Models for Security. GraMSec 2017. Lecture Notes in Computer Science, vol 10744. Springer, Cham

Visualizing Cyber Security Risks with Bow-Tie Diagrams

Karin Bernsmed[1], Christian Frøystad[1], Per Håkon Meland[1,3], Dag Atle Nesheim[2], and Ørnulf Jan Rødseth[2]

[1] SINTEF Digital
{karin.bernsmed,christian.froystad,per.h.meland}@sintef.no
[2] SINTEF Ocean
{dag.atle.nesheim,ornulfjan.rodseth}@sintef.no
[3] Norwegian University of Science and Technology
per.hakon.meland@ntnu.no

Abstract. Safety and security risks are usually analyzed independently, by different people using different tools. Consequently, the system analyst may fail to realize cyber attacks as a contributing factor to safety impacts or, on the contrary, design overly secure systems that will compromise the performance of critical operations. This paper presents a methodology for visualizing and assessing security risks by means of bow-tie diagrams, which are commonly used within safety assessments. We outline how malicious activities, random failures, security countermeasures and safety barriers can be visualized using a common graphical notation and propose a method for quantifying risks based on threat likelihood and consequence severity. The methodology is demonstrated using a case study from maritime communication. Our main conclusion is that adding security concepts to the bow-ties is a promising approach, since this is a notation that high-risk industries are already familiar with. However, their advantage as easy-to-grasp visual models should be maintained, hence complexity needs to be kept low.

Keywords: security, safety, risk assessment, bow-tie diagrams, maritime communication

1 Introduction

One of the least understood challenges for cyber physical systems (CFS) is uncertainty in the environment, cyber attacks and errors in connected physical devices [46]. The tight coupling between the cyber and physical world leads to new forms of risks that have not been considered adequately, such that the cyber element adversely affects the physical environment [4]. Safety risks, where the system can harm the environment in which it operates, and security risks, where the environment (e.g. malicious actors and other systems) can harm the system, tend to be analyzed independently [42], by different people using different standards, tools and notations. As pointed out by Sun et al. [50], safety and security goals interact synergistically or conflictingly, and should therefore be evaluated together. If not, conflicts can result in either (a) overly secure systems that compromise the reliability of critical operations or (b) create insecure systems where back-doors are easily found.

An inherent challenge when combining safety and security in an analysis is the increased complexity. Graphical visualizations are helpful when you want to make complex problems easier to understand and navigate [20]. The purpose of this paper is to bridge the gap between safety and security during risk assessment by utilizing the graphical bow-tie diagram methodology [14, 11, 15, 25]. Bow-tie diagrams are very suitable for communicating the results of a risk assessment to different stakeholders within an organization due to the clear diversification of causes and effects for a given unwanted event, and to clarify which barriers have (or have not) been implemented. Bow-tie analysis, which includes the generation of one or more bow-tie diagrams, is a common approach to map the risks associated with unwanted events in, for example, the oil and gas industry. Our approach is to take advantage of the familiarity of this graphical notation among industry experts, analyze use cases within the safety-critical maritime sector, and try to answer the following research questions:

1. How can bow-tie diagrams be extended to include security considerations in addition to safety considerations?
2. How can the likelihood of cause and severity of cyber attacks be visualized in bow-tie diagrams?

In order to answer these questions, we apply a *design science* research methodology [48], with focus on the extended bow-tie diagram methodology as an artefact with a high priority on relevance for the cyber physical domain. Evaluation is done through analysis of descriptive, constructed use cases for maritime service scenarios to demonstrate its utility [21].

Our goal has not been to create yet another theoretical model for risk assessment, but to propose a solution to a real, existing problem we experience in the maritime domain when introducing new technology that may have effect both safety and security. This follows the research paradigm of *pragmatism* [19], which is associated with action, intervention and constructive knowledge. Furthermore, it should be based on real problems and have practical usefulness beyond the specific case studies.

This paper is organized as follows. Section 2 presents related work. In Section 3, we introduce the marine communication case study in which we have developed the proposed methodology. Section 4 explains the concepts and terminology that we use and Section 5 presents the proposed bow-tie risk assessment methodology, which is exemplified in Section 6. Finally, in Section 7 we discuss the results and Section 8 concludes the paper.

2 Related work

The most common way of documenting and visualizing risks is in a risk matrix, where the seriousness of the evaluated risks can be easily compared based on the combination of likelihood and consequence. The US Air Force developed the Risk Matrix Approach (RMA) [18] in 1995, and after that it has spread out to a multitude of domains, such as weapons manufacturing, finance, transport and project management [38]. Still, RMA is a very simplistic notation that does not properly visualize the causes of the risks, and how to address them.

Within the field of security, there are many more specialized modelling notations that are in general concerned about *"identifying system behavior, including any security defenses; the system adversary's power; and the properties that constitute system security"* [5]. Security modelling comes in many different forms and flavors, but they all share the common aim of understanding security issues so they can be dealt with effectively. Which one to choose usually depends on what the analyst wants to focus on, level of abstraction/details and personal preference (e.g. familiarity). To quote Shostack [47]: *"different diagrams will help in different circumstances"*. For instance, an attack tree [45, 31] is a tree-based notation showing how an adversary can choose among different paths or branches to obtain an overall attack goal. The attack-defense trees [26] extend this notation by also adding preventive nodes, which again can be attacked by attack nodes. Attack graphs [40] and vulnerability cause graphs [8] are examples of a graph-based notation used for analyzing vulnerabilities, and CORAS [30] contains several graphical notations for a risk analysis process. There also exist different types of security extensions to more general purpose graphical modelling notations, such as Data flow diagrams [47], UML [24, 49] and BPMN [32].

For safety, there are many notations that go even further back in history. The fault-tree analysis (FTA) method was developed in the 1960s for safety and reliability [29], and a recent survey of usage is provided by Ruijters and Stoelinga [43]. Event tree analysis (ETA) is an established technique originating from the nuclear industry [3], and is used to analyze how a series of events can lead to a potential accident scenario. Similarly to ETA, cause-consequence diagrams (CCA) [39] are also used to analyze safety causes.

When considering safety and security in combination, there have been quite a few related studies. For instance, Winther et al. [52] show how to handle security issues as part of HAZOP studies, which is a systematic analysis on how deviations from the design specifications in a system can arise, and whether these deviations can result in hazards. Raspotnig et al. [42] have use UML-based models within a combined safety and security assessment process to elicitate requirements. Bieber and Brunel [7] show how common system models for security and safety can be used for airworthiness certification within aviation. Kumar and Stoelinga [28] have married fault and attack trees so that both safety and security can be considered in combination. Further examples of methods, models, tools and techniques in the intersection of safety and security can be found in the surveys by Zalewski et al. [53], Piètre-Cambacédès and Bouissou [41], Chockalingam et al. [12], as well as Kriaa et al. [27].

There have been several efforts by practitioners related to the use of bow-tie diagrams for cyber security, but they differ from what we are presenting in this paper in several ways. For instance, a report from SANS Institute [35] outlines how a bow-tie risk assessment methodology can be applied to conduct a cyber security risk assessment in an engineering environment. There is no change to the diagram notation as such, but they argue that *"the first step towards obtaining Engineering community buy-in"* is to compare concepts from security to bow-tie, and basically evaluate cyber threats in the same manner as hazards. They also include considerations related to actors and motivation, but this is done in order to reduce the number of possible scenarios before modelling, and not part of the notation itself. A report from DNV-GL [16] also proposes the use of bow-tie diagrams as a key component in a cyber security assessment program for the maritime sector. Here, standard safety notation is used, and the focus is on visualization of barriers. Quantitative indicators are explicitly left out, and even though vulnerability consideration is central in the overall assessment process, this is not included as diagram concepts. Similarly, the *Bow Tie for Cyber Security* series [22] at PI Square gives numerous examples where the standard notation is used for security. The US Coastguard has also published a report [34] on how to use bow-ties to identify preventive and responsive responses to cyber attacks for marine transportation systems. Their examples are on a very high abstraction level, where causes are for instance *hactivists*, *technical errors* and *insider threats*. Two additional examples of bow-tie diagrams that visualize IT security risks are provided in [10]. The focus here is more on chains of barriers, although it seems like vulnerabilities are represented as escalation factors.

3 Case study: Maritime communication

In order to give a better understanding of the methodology and examples used in the later sections, we would like to explain our maritime case study and why security is a growing concern intertwined with safety in this domain.

Shipping has become increasingly dependent on digital data exchanges. As dependence grows and the functions supported becomes more entangled in the ship operations and critical interactions with on-shore authorities, the need to consider consequences of digital attacks on the data exchanges also increases. This calls for a more systematic approach to maritime cyber security.

In 2011, ENISA pointed out [13] that the *"awareness on cyber security needs and challenges in the maritime sector is currently low to non-existent"*. Come 2015, the Lysne commission of Norway [2] reaffirmed this message. The lack of general awareness regarding cyber security, makes the industry more vulnerable to attacks.

Maritime navigational systems of today rely heavily on Global Navigation Satellite Systems (GNSS), such as GPS and GLONASS, to navigate safely, avoid collisions or groundings and for voyage optimization. The GNSS signals available for civilians are unencrypted and unauthenticated and are easily jammed or even spoofed [6]. Automatic Identification System (AIS) is used to identify other ships and their intentions, but can also be used to transmit short safety messages, e.g. to act as virtual aids to navigations. AIS is becoming part of the more extensive VHF Data Exchange System (VDES), which will extend the use of AIS to include even more digital information exchanges. The AIS messages are unencrypted and unauthenticated, and relatively easy to jam or spoof. Furthermore, IOActive [44] conducted tests on SATCOM firmware from multiple vendors and found vulnerabilities such as hardcoded credentials, undocumented protocols, insecure protocols, backdoors, and weak password reset. Our attention is on digital data exchanges between ships and between ship and shore and the possible consequences of cyber-attacks on these exchanges.

Ships spend most of their time at sea with a minimal crew, and remote monitoring and maintenance is becoming more and more common. If not organized in an appropriate way, this could allow an attacker extensive and easy access to the systems on the ship. Additionally, there are multiple actors connected to the network on-board a ship, including passengers, crew, and operational systems. These actors have different requirements regarding safety, security and separation. For instance, some vessels have physically separated networks, while others only provide logically separated networks. The mechanisms for logical separation of networks vary, but are often just a simple firewall.

4 Concepts and terminology of bow-ties

A bow-tie diagram is shaped like a bow-tie, where the central *knot* typically represent an accident scenario, or as we will later refer to, an unwanted event. The diagram can be seen as a combination of a fault tree and an event tree [17], where the left side shows which causes can lead up to the accident, and the right side the potential effects once the accident has occurred. As pointed out by the tool provider CGE Risk Management[4], the power of this diagram is that it gives a clear distinction between proactive and reactive risk management, in combination with an overview of multiple plausible scenarios.

To combine security with bow-tie safety assessment, we need to synchronize the terminology and concepts from the safety and security domains. The bow-tie diagram in Fig. 1 shows the traditional layout, notation and concepts from safety assessments in the upper left horizontal part (*cause, barrier, escalation factor*), with concepts we introduce from security in the lower left horizontal part (*threat, security control*). *Hazard* and *unwanted event* are mainly from safety, while *asset* comes from security. On the right side of the figure, the *consequence* concept is shared between safety and security, and can be remedied with safety barriers and security controls, often in combination. We describe these concepts further below.

As defined by International Maritime Organization (IMO) [23], the first step in a Formal Safety Assessment (FSA) [23] is to identify all potential hazards that can contribute to accidents. A *hazard* is a potential to threaten human life, health, property or the environment. Examples of maritime hazards are off-shore operations, hazardous substances and sources of ignition onboard and external hazards, such as storms, lightening and other ships. Hazards may give rise to

[4] https://www.cgerisk.com/knowledge-base/risk-assessment/thebowtiemethod

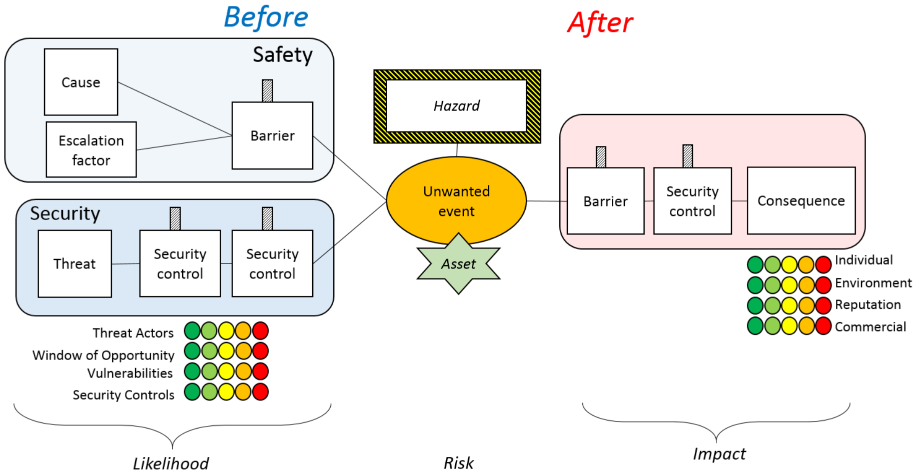

Fig. 1. Our combined approach for modelling safety and security in a bow-tie diagram.

scenarios in which people, the environment or property will be damaged. The list of identified hazards and their associated scenarios will be used as input to the safety risk assessment. Basically, a hazard can be anything with the potential to cause harm, but which is also necessary to perform business. From a risk analysis perspective, the hazard needs to be controlled so that unwanted events will not occur.

An *unwanted event* in safety assessment, also known as top event, loss event, or loss of control, represents what will happen if one loses control over a hazard, which again can have severe *consequences*. An unwanted event is typically caused by an accident, or a random failure. In security assessments, the equivalent is often called *incident*, something that typically affects the confidentiality, integrity or availability of a critical system, data, or processes necessary for the operation of the business. Such incident may have malicious or accidental causes. In our model, we are using the term unwanted event for anything that can cause harm to the asset(s) associated to the hazard, regardless if they stem from safety or security causes. In real life, it is often a combination of different causes that lead to unwanted events, therefore we want to evaluate them together.

Related to security, an *asset* is anything that has value to an organization. The ISO/IEC 27005 standard [1] distinguishes between primary assets, which are core business processes and their corresponding information, whilst supporting assets are those required to be in place to support the activities of the primary assets. Typical examples of (primary) assets in a maritime context are Maritime Safety Information (MSI), ship certificates, and electronic nautical charts. Asset is not a concept that is used in traditional safety assessment, but is usually the first thing to identify when it comes to security assessments. Therefore, we include a mapping between hazard and which assets will be damaged in case the unwanted event occurs.

A *threat* is anything that can potentially cause an unwanted event [1]. Within safety assessments, the term *cause* is very often used directly for the same meaning. A *barrier* is a mechanism that aims to interrupt causes of unwanted events, or that it is possible to recover from the unwanted event without severe consequences. In a security context, the term barrier corresponds to the term *control*, which is a means of managing risk, including policies, procedures, guidelines, practices or organizational structures, which can be of administrative, technical, management, or legal nature [1]. These can be preventive controls used to avoid, detect or mitigate threats, or reactive controls, which are intended to limit the damage caused by an incident. Note that in a security context, the word safeguard, mitigations, or countermeasure, are sometimes used as a synonym for control. An *escalation factor* is anything that may cause a safety barrier to fail. There is no one-to-one mapping between this concept and security terminology, however, to succeed with a threat, a threat actor will need to exploit one or more vulnerabilities, which often is only feasible at a certain point of time (window of opportunity).

In our model, we use threats to explicitly represent malicious activities, while causes are related to traditional safety accidents. We continue to use both barrier and security control for both sides of the bow-tie, though they may have the same implementation (e.g. through redundancy). Note that there can be chains of both barriers or security controls (the latter is illustrated in Fig. 1). Such chains follow the principle of *defence in depth* - if the first barrier fails or control is circumvented, there is another one still operating.

We also introduce a set of color coded indicators for each threat branch on the left side, and for each consequence branch on the right side of the diagram. These indicators are meant to help visualize the likelihood of an unwanted event, and the severity of a consequence in similar manner that is used for risk matrices. This allows us to adopt the RMA framework as described in Section 2 as apart of the notation, and make use of the color indicators that the industry community is already familiar with. For a threat branch, we associate indicators related to *threat actors, window of opportunity, vulnerabilities* and *security controls.* For instance, the threat actors indicator informs whether or not it is likely that there exists groups or individuals who have the competence, resources and motivation necessary to perform an attack and instantiate the threat. Similarly, we indicate the likely existence of the other indicators. For a consequence branch, the indicators represent the severity of the impact related to *individuals*, the *environment*, the *reputation* of a company and *commercial* (monetary) loss.

In the next section, we focus on how to identify what color should be used for each indicator, and how to quantify the overall risk of a bow-tie diagram for an unwanted event.

5 Risk assessment

As illustrated in Fig. 1, the risk of an unwanted event will be a combination of the likelihood and the impact of the unwanted event. Our contribution in this paper focuses on a subset of all potential unwanted events, which are those caused by hostile cyber attacks. In our model, an unwanted event U will be a function of one or more threats. Each unwanted event will lead to one or more consequences C, where each identified consequence is associated with a corresponding impact (i.e. severity, or loss,) value L. The risk R associated with a certain unwanted event U, which we denote $R(U)$, will then be approximated as the probability that the unwanted event occurs, i.e. $p\,(U)$, multiplied with the worst-case consequence impact value that has been identified, which we denote L_C, and the likelihood that this consequence occurs, i.e. $p\,(C)$. The formal expression for this is

$$R(U) \approx p\,(U) \ \times L_C \ \times p\,(C) \tag{1}$$

To quantify the risk of an unwanted event, we hence need to assess 1) the probability of the unwanted event (as a function of one or more identified threats) and 2) the impact value and probability of the worst-case consequence of the unwanted event.

5.1 Assessing the left side of the bow-tie (cause)

Assessing the probability of a cyber attack is a notoriously difficult problem. In our model, we assume that all the threats are *mutually independent*. This means that all the identified cyber attacks will be executed independently of each other and that any of them can manifest itself and cause the unwanted event during the time for which the system, or service, is being assessed. Under this assumption, the probability of the unwanted event U can be computed as

$$p\,(U) = p\,(at\ least\ one\ T_i\ occurs) = 1 - \prod_{i=1}^{n} (1 - p\,(T_i)) \tag{2}$$

where $p(T_i),\ i = 1 \ldots n$, is the probability of threat T_i. The problem will hence be reduced to assessing the probabilities, or likelihoods, of the individual threats that have been identified.

Compared to more simplistic probability models, in which the threats are modelled as mutually exclusive (i.e. $p\,(U)$ will be computed as a sum of the individual threats), the proposed Equation 2 is much more realistic, since it allows more threats to manifest within the same time interval, which corresponds more closely to the real world. By using Equation 2, we can also model cases in which multiple attackers work simultaneously to exploit different vulnerabilities, and cases where one attacker exploits all the vulnerabilities he can find. However, the assumption that all the threats are independent may not always be true. In particular, it is questionable whether one can model scenarios in which an attacker is aware of all the potential threats that can be carried out, since this may affect the probabilities of the individual threats, hence violating the independence assumption. Another issue may be that, for some unwanted events, once the unwanted event has happened, it will be less likely to happen again due to increased awareness. This is a common situation in an security context, where threats are manifesting themselves through the actions of human beings rather than through random failures, and the malicious actors will lose their *element of surprise.*

Another characteristic of Equation 2 is that the more threats one identifies, the higher the probability of the unwanted event. A side effect of using this model could therefore be that a more thorough risk assessor, who manages to identify more threats, will also end up with a higher probability of the unwanted event. However, the influence of the number of identified threats will be negligible, as long as both the threat probabilities and the number of identified threats are sufficiently small (which is the case in most real-life scenarios).

In our opinion, in spite of the aforementioned issues, this is the simplest and most straightforward alternative we have for computing the probability of an unwanted event $p\left(U\right)$ as a function of the identified threats. This same model is frequently used in system reliability analysis, in which a system analysist models the system as a set of components, assesses the individual failure rates of the components and evaluates the effect of the total system reliability. In our case, we model malicious threats rather than random failures, however, the underlying line of thought is similar; we are considering multiple sources of error that can cause the system, or service, to fail, regardless of cause. Note that, when using this approach, care must be taken to ensure that all the identified threats are independent and, as explained above, the risk assessor must understand the characteristics of the underlying mathematical model.

Assessing the threat actors, window of opportunity, vulnerabilities and security countermeasures. We move on to describe how factors, such as the actors who pose the threat, the needed window of opportunity for the threat to be successful and any vulnerabilities and security countermeasures present in the system can be assessed and visualized. As explained in Section 4, we use color coded indicators to represent these factors in the graphical model.

Threat actors Threat actors are the attackers who will represent a security risk against the system that is being assessed. Threat actors can be classified in terms of characteristics, such as skill, capabilities, resources, intent and access [9]. The risk assessor can estimate the threat actors by using the values of Table 1.

Window of Opportunity The "window of opportunity" depends on how often/long the threat actor theoretically could gain access to the target (system or data) and how often/long the target of interest is within reach of the attacker. The risk assessor can estimate the window of opportunity by using Table 2.

Vulnerabilities No system is perfect, nor are the security measures that are put in place to prevent the threat from manifesting itself. Vulnerabilities can range

Table 1. Color coding for representing the threat actors

Threat actors

Dangerousness	Description	Color coding
Severe	There are threat actors highly capable of pursuing this threat	🟥
High	There are threat actors capable of pursuing this threat	🟧
Moderate	There are threat actors somewhat capable of pursuing this threat	🟨
Low	There are threat actors interested in pursuing this threat, but their capability is limited	🟩
None	There are threat actors interested in pursuing this threat, but they are not capable of acting on this interest	🟩

Table 2. Color coding for representing the window of opportunity

Window of opportunity

Window	Description	Color coding
Always	This threat is always possible.	🟥
Frequent	This threat is frequently possible (there will be an opportunity about once every week).	🟧
Rare	This threat is rarely possible (there will be an opportunity about once every year).	🟨
Extremely rare	This threat is extremely rarely possible (there will be an opportunity about once every 10th year).	🟩
Never	This threat is never possible.	🟩

from simple programming errors to large design flaws of software, hardware and processes. The presence of vulnerabilities increases the likelihood of a threat manifesting. The risk assessor can estimate the existence of vulnerabilities by using Table 3.

Table 3. Color coding for representing the presence of vulnerabilities

Vulnerabilities

Vulnerability	Description	Color coding
Known easy	One or more known vulnerabilities exist, which are easy to exploit.	🟥
Known-difficult	One or more known vulnerabilities exist, but they are either not publicly known, or they are difficult to exploit.	🟧
Unknown	No known vulnerabilities exist, however, vulnerabilities are expected to appear in the near future.	🟨
Very unlikely	It is very unlikely that the system has, or will have, any vulnerabilities in the near future.	🟩
Formally proven absence	Formal methods, or the like, have been applied to demonstrate that no vulnerabilities exist. It is extremely unlike that vulnerabilities will appear in the near future.	🟩

Security controls Finally, the risk assessor will need to input information about the existence of security control and assess their effectiveness (Table 4).

Assessing the threats For each threat T_i and preventive security controls $Ctrl_1 \ldots Ctrl_m$, the risk assessor choose values for *Threat Actors*, *Window of Opportunity*, *Vulnerabilities* and *Security Controls* according to Table 1, 2, 3 and 4. This is visualized as extended traffic lights as shown in Figure 2. In addition to the traffic lights, the relevant controls for each threat are shown as separate boxes to give an overview of which threats are mitigated by which controls.

The visualization in Fig. 2 serves as domain specific assistance to the risk assessor when assessing $p(T_i),\ i = 1 \ldots n$, i.e. the probability of each of the identified threats. We do not dictate exactly how this estimation should be done in practice, as there are different ways of doing threat prediction, and any model depends a lot on the available information used as input. When working with maritime threat scenarios, we have been using averages from generic threat intelligence data, and then adjusted these based on the case specific domain data using expert opinions.

5.2 Assessing the right side of the bow-tie (consequence)

The consequence of an evaluated risk can manifest itself in many ways. FSA normally only consider individual risk and societal risk which represents the main

Table 4. Color coding for representing the effectiveness of security controls

Security controls

Control	Description	Color coding
Known to be ineffective	No security countermeasure exists, or, one or more security countermeasures exists but they are known to be ineffective.	🟥
Probably not effective	One or more security countermeasures exists but they can be circumvented.	🟧
Effective	One or more security countermeasures exists, which are believed to be effective.	🟨
Very effective	One or more security countermeasures exists, which are very effective.	🟩
Formally proven effective	Formal methods, or the like, have been applied to demonstrate that existing security mechanisms are sufficient and work as intended.	🟩

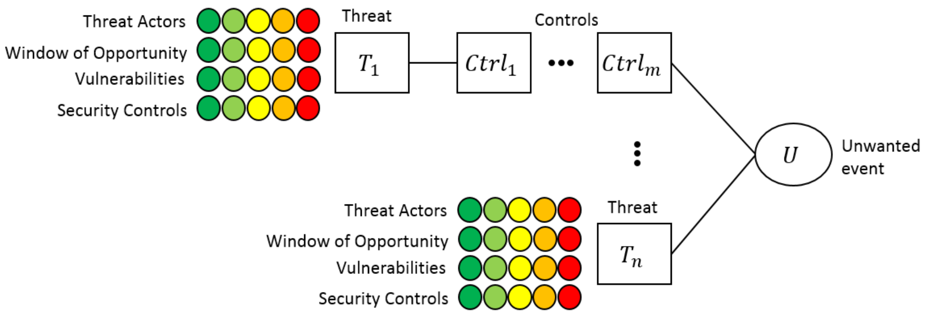

Fig. 2. The relation between an unwanted event, threats, threat actors, window of opportunity, vulnerabilities and (preventive) security controls

scope of the Maritime Safety Committee in IMO where the FSA was developed. We have found it useful to also include other aspects, such as the environmental (pollution), commercial (monetary losses) and reputational (loss of confidence by e.g. customers, business partners, bank, insurance, regulatory bodies) damage caused by each identified unwanted event in our model. As an example of reputational damage, the Paris MoU[5] publishes a black list for all ships depending on results from Port State Controls. Once your ship is on this list, you are much more eligible for inspections and your operation may suffer.

Table 5. Consequence type and severity level

Consequences

Level	Individual	Environment	Reputation	Commercial	Color coding
Catastrophic	Multiple deaths	Uncontained release with potential for very large environmental impact	International coverage, unrecoverable damage	$ 50 000 k	🟥
Critical	One death	Uncontained release with potential for major environmental impact	National and some international coverage, impact lasting more than a year	$ 5 000 k	🟧
Moderate	Multiple severe injuries	Uncontained release with potential for moderate environmental impact	National media coverage, impact lasting more than 3 months	$ 500 k	🟨
Negligible	One minor injury	On site release contained without external assistance	Local complaint/ recognition, impact less than one month	$ 5 k	🟩
None	No injuries	No effect	No damage	$ 1 k or less	🟩

The risk assessor can estimate the consequence of each identified unwanted event using Table 5. One obvious problem with comparing these different outcomes is to compare consequences for life and health with purely economic or environmental damages. However, it is possible to compare the economic consequences of a lost life or health damage to other more direct economic consequences of a cyber attack. Our approach is to follow this (semi-) quantitative assessment, and leave a more qualitative societal risk acceptance analysis to later stages.

Individual consequence represents the direct danger to life or health of persons on board the ship, on other ships or on shore. It does not include secondary effects due to, e.g. pollution or other factors. As noted above, it is not trivial

[5] https://www.parismou.org/

to assess the value of life and health in purely economic terms. The problem is, for instance, complicated by the different economic values assigned to lives in different parts of the world [51]. For example, this value was estimated to be at USD 0.8 million in South Korea in the year 2000, and at USD 9.7 million in Japan the same year. In our model, we will use the mean value of USD 5 million for one life as baseline. This represents the mean value from [51], but not weighted according to population in the different areas.

We follow the defined severity levels for economical loss as shown in Table 5. This maps critical to the above value corresponding to loss of one life and adjusts other levels accordingly.

The inclusion of reputational and economical loss in the risk assessment has been a matter of some discussion. Our rationale for doing this and not only focusing on individual and environmental risks, is that in many cases the motivation for and the consequences of a successful cyber-attack is likely to be much higher in the commercial domain than in the general safety domains. This assumption is strengthened by todays ship bridge operational regime where all received information must be checked against other sources of information, including making visual assessment of the ships situation. Thus, including commercial consequences will likely lead to more risks being assessed as not acceptable and by that lead to a higher overall safety level.

6 Use case example: Navigational Information Update

In this section, we demonstrate the use of our proposed methodology to represent unwanted events in a bow-tie diagram and to assess the corresponding risk. The context is cyber security threats in the maritime communication case study introduced in Section 3. The use case we investigate is called *Navigational Information Update*. The objective here is to illustrate the visualization, and not to present the complete description.

Ships are required to keep critical electronic databases up to date. Such databases include electronic charts and lists of navigation signals. Updating can be done by requesting updates as the voyage progresses and getting data from the chart provider. In the near future, this will be implemented over an Internet based service via satellite or other high capacity carriers. Failing to get the right data can cause safety hazards as well as a danger of detention by the Port State Control in the next port. In addition, some of this information is provided by commercial companies that need to protect the supplied information from copying to non-paying ships.

In this example, we address electronic ship navigation as a potential hazard and we want to assess the risk of the unwanted event “Ship receives incorrect updates”. The affected asset is the navigation data that is being transferred. Fig. 3 and 4 illustrate the identified threats, security controls and potential consequences that we have identified in our analysis.

To compute the risk, we need to assess the probabilities of all the identified threats, as well as the impact value and probability of the worst-case consequence

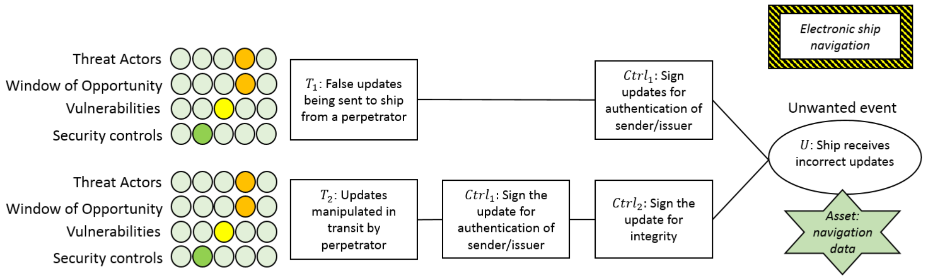
Threat Actors
Window of Opportunity
Vulnerabilities
Security controls
T_1: False updates being sent to ship from a perpetrator
$Ctrl_1$: Sign updates for authentication of sender/issuer
Electronic ship navigation
Unwanted event
U: Ship receives incorrect updates
Asset: navigation data
Threat Actors
Window of Opportunity
Vulnerabilities
Security controls
T_2: Updates manipulated in transit by perpetrator
$Ctrl_1$: Sign the update for authentication of sender/issuer
$Ctrl_2$: Sign the update for integrity

Fig. 3. The left hand threat side with preventive controls for the unwanted event "Ship receives incorrect updates"

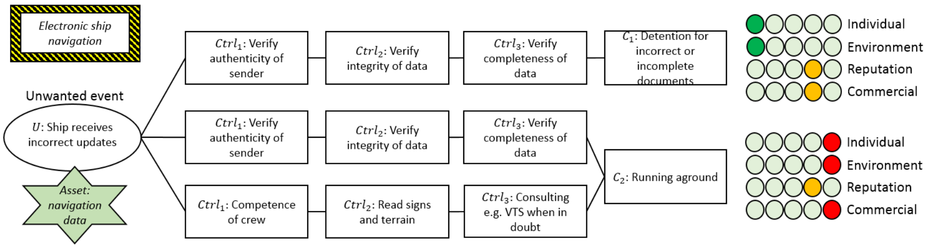
Electronic ship navigation
Unwanted event
U: Ship receives incorrect updates
Asset: navigation data
$Ctrl_1$: Verify authenticity of sender
$Ctrl_2$: Verify integrity of data
$Ctrl_3$: Verify completeness of data
C_1: Detention for incorrect or incomplete documents
Individual
Environment
Reputation
Commercial
$Ctrl_1$: Verify authenticity of sender
$Ctrl_2$: Verify integrity of data
$Ctrl_3$: Verify completeness of data
C_2: Running aground
Individual
Environment
Reputation
Commercial
$Ctrl_1$: Competence of crew
$Ctrl_2$: Read signs and terrain
$Ctrl_3$: Consulting e.g. VTS when in doubt

Fig. 4. The right hand consequence side with reactive controls for the unwanted event "Ship receives incorrect updates"

identified for this unwanted event. The assessment of a risk assessor, who has considered the threat actors, window of opportunity, vulnerabilities and security controls, is used as a source for this threat prediction. If we for instance set probability of threat $T_1 = 0.45$ and probability of threat $T_2 = 0.23$, and then apply Equation 2, we can compute the probability of the unwanted event:

$$p(U) = 1 - (1 - p(T_1)) \times (1 - p(T_2)) = 1 - (1 - 0.45) \times (1 - 0.23) \approx 0.57 \quad (3)$$

Furthermore, let's assume the consequence $C_1 = 0.3$, $p(C_1) = 0.5$, $C_2 = 0.7$ and that $p(C_2) = 0.2$. By applying Equation 1, we find that the risk of the unwanted event to be:

$$R(U) \approx 0.57 \times 0.7 \times 0.2 \approx 0.08 \quad (4)$$

This number does not mean much by itself, but can be used as a relative number when comparing with other unwanted events, and to justify the addition of barriers/controls.

As illustrated by this simple example, the bow-tie diagram provides an illustrative overview over the identified threats, security controls and potential consequences of the unwanted event.

7 Discussion

To make useful cyber security visualizations with bow-tie diagrams, we needed to identify which security concepts to include and what kind of quantified input data would be meaningful as input to the diagrams. In our case, we have done this in separate processes, one for each side of the diagrams. For the left side (potential causes and threats, including likelihood), security and domain experts participated in a workshop setting (n=10), while the right side (consequences and their severity) was evaluated by representatives from maritime industry and coastal authorities through an online survey (n=18). Both groups were working with the same set of seventeen service scenarios for maritime communication, and twenty use cases that overlapped between the services. Note that none of these groups worked directly with bow-ties as a graphical notation, but were focused on types of threats, consequences and estimating values based on their experience and expert opinion. Based on these results, which are documented in [36], we have developed the methodology for visualizing concepts and quantified values for cyber security with the bow-tie notation, addressing research question 1 from Section 1. This has then been applied to a sample of the use cases from the service scenarios, as shown in Section 6, to demonstrate the utility of our approach. We consider this to be a first step of evaluation, where we have shown that the main security concepts can be contained and visualized. We have also tried to address research question 2 by adding color coded indicators to the diagrams, which are there to justify the likelihood and impact of an unwanted event. However, further work is needed to do in-depth evaluation on how this is perceived and found useful by other analysts, stakeholders from the maritime domain, as well as stakeholders from other safety domains.

Some general observations we have made when working with bow-tie modelling is that they are very suitable to show the broadness and distribution of different causes and consequences for unwanted events, along with protective and reactive barriers. However, this approach also has its limitations. For instance, a bow-tie diagram will struggle to represent the depth and details of how attacks can be performed. Furthermore, a single cause or threat can lead to different unwanted events, therefore, there can easily be repetition/redundancy between a collection of bow-ties addressing different hazards. We therefore recommend that the diagrams are complemented with more established methods for threat modelling, and that these are reused and referred to from nodes within the bow-ties. This can for instance be fault-trees for safety, or generic attack trees or misuse cases for security, that Meland et al. [33] have already showed can be shared and reused between different projects, organizations or domains with benefit. A prerequisite to realize this would be modelling tool support beyond simple drawing tools, as well as collaboration and willingness to share knowledge between risk analyst addressing both safety and security.

To capture more security related information within a bow-tie, it is also possible to add specific nodes in the model for concepts such as threat actors and vulnerabilities. We believe that this would lead to an unnecessary complexity of the diagram, and it would lose some of its advantage as an easy to grasp graphical representation. The number and types of nodes would increase, and there would in many cases be many-to-many relationships between threat actors, threats, vulnerabilities, and security controls. Therefore, we rather use the more simplified notation of indicators related threat and consequence branches, that sums up for instance whether it is likely there are many relevant threat actors.

8 Conclusion

Safety assessments with bow-tie diagrams give a good pictorial understanding of major risks and how they are controlled. This is a technique that many of the high-risk industries are already familiar with, such as oil and gas, mining, aviation, maritime and public health services [37]. Due to the increasing connectivity of cyber physical systems, these are the same industries that are now becoming more and more exposed to cyber attacks. To avoid conflicting goals and requirements between safety and security, we believe that adding security to the bow-tie notation is more accommodating than inducing yet another specialized, separate modelling technique that tries to capture all aspects of safety and security. Bow-tie diagrams are meant to be easy to understand, and by combining a minimal set of security concepts along with associated indicators, we can show both safety and security considerations without overflowing the diagrams.

Acknowledgments The research leading to these results has been performed as a part of the Cyber Security in Merchant Shipping (CySiMS) project, which received funding from the Research Council of Norway under Grant No. 256508, and the SafeCOP-project, which received funding from the ECSEL Joint Undertaking under Grant No. 692529. We appreciate all the feedback and comments from Professor Guttorm Sindre at NTNU and anonymous reviewers that helped us improve this paper.

References

1. ISO/IEC 27005 Information technology – Security techniques – Information security risk management. Tech. rep. (2008), http://www.iso.org/iso/catalogue_detail?csnumber=56742
2. Digitale Sarbarheter Maritim Sektor. Tech. rep. (2015), https://www.regjeringen.no/contentassets/fe88e9ea8a354bd1b63bc0022469f644/no/sved/7.pdf
3. Andrews, J.D., Moss, T.R.: Reliability and risk assessment. Wiley-Blackwell (2002)
4. Banerjee, A., Venkatasubramanian, K.K., Mukherjee, T., Gupta, S.K.S.: Ensuring safety, security, and sustainability of mission-critical cyber–physical systems. Proceedings of the IEEE 100(1), 283–299 (2012)
5. Bau, J., Mitchell, J.C.: Security modeling and analysis. IEEE Security & Privacy 9(3), 18–25 (2011)
6. Bhatti, J., Humphreys, T.: Hostile control of ships via false gps signals: Demonstration and detection. Navigation (2016)

7. Bieber, P., Brunel, J.: From safety models to security models: preliminary lessons learnt. In: International Conference on Computer Safety, Reliability, and Security. pp. 269–281. Springer (2014)
8. Byers, D., Ardi, S., Shahmehri, N., Duma, C.: Modeling software vulnerabilities with vulnerability cause graphs. In: Proceedings of the International Conference on Software Maintenance (ICSM06). pp. 411–422 (2006)
9. Casey, T.: Threat agent library helps identify information security risks (2007), https://communities.intel.com/docs/DOC-1151
10. CGE Risk Management Solutions: Using bowties for it security (2017), https://www.cgerisk.com/knowledge-base/risk-assessment/using-bowties-for-it-security
11. Chevreau, F.R., Wybo, J.L., Cauchois, D.: Organizing learning processes on risks by using the bow-tie representation. Journal of hazardous materials 130(3), 276–283 (2006)
12. Chockalingam, S., Hadziosmanovic, D., Pieters, W., Teixeira, A., van Gelder, P.: Integrated safety and security risk assessment methods: A survey of key characteristics and applications. arXiv preprint arXiv:1707.02140 (2017)
13. Cimpean, D., Meire, J., Bouckaert, V., Vande Casteele, S., Pelle, A., Hellebooge, L.: Analysis of cyber security aspects in the maritime sector (2011)
14. Cockshott, J.: Probability bow-ties: a transparent risk management tool. Process Safety and Environmental Protection 83(4), 307–316 (2005)
15. De Dianous, V., Fiévez, C.: Aramis project: A more explicit demonstration of risk control through the use of bow–tie diagrams and the evaluation of safety barrier performance. Journal of Hazardous Materials 130(3), 220–233 (2006)
16. DNV-GL AS: Recommended practice. cyber security resilience management for ships and mobile offshore units in operation (2016)
17. Ferdous, R., Khan, F., Sadiq, R., Amyotte, P., Veitch, B.: Analyzing system safety and risks under uncertainty using a bow-tie diagram: An innovative approach. Process Safety and Environmental Protection 91(1), 1–18 (2013)
18. Garvey, P.R., Lansdowne, Z.F.: Risk matrix: an approach for identifying, assessing, and ranking program risks. Air Force Journal of Logistics 22(1), 18–21 (1998)
19. Goldkuhl, G.: Pragmatism vs interpretivism in qualitative information systems research. European Journal of Information Systems 21(2), 135–146 (2012)
20. Hall, P., Heath, C., Coles-Kemp, L.: Critical visualization: a case for rethinking how we visualize risk and security. Journal of cybersecurity 1(1), 93–108 (2015)
21. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. MIS Q. 28(1), 75–105 (Mar 2004), http://dl.acm.org/citation.cfm?id=2017212.2017217
22. Hpaul: Security: Bow Tie for Cyber Security (0x01): Ho... — PI Square (2016), https://pisquare.osisoft.com/groups/security/blog/2016/08/02/bow-tie-for-cyber-security-0x01-how-to-tie-a-cyber-bow-tie
23. IMO: Revised guidelines for formal safety assessment (fsa) for use in the imo rule-making process (2013)
24. Jürjens, J.: Umlsec: Extending uml for secure systems development. In: International Conference on The Unified Modeling Language. pp. 412–425. Springer (2002)
25. Khakzad, N., Khan, F., Amyotte, P.: Dynamic risk analysis using bow-tie approach. Reliability Engineering & System Safety 104, 36–44 (2012)
26. Kordy, B., Mauw, S., Radomirović, S., Schweitzer, P.: Foundations of attack–defense trees. In: International Workshop on Formal Aspects in Security and Trust. pp. 80–95. Springer (2010)

27. Kriaa, S., Pietre-Cambacedes, L., Bouissou, M., Halgand, Y.: A survey of approaches combining safety and security for industrial control systems. Reliability Engineering & System Safety 139, 156–178 (2015)
28. Kumar, R., Stoelinga, M.: Quantitative security and safety analysis with attack-fault trees. In: High Assurance Systems Engineering (HASE), 2017 IEEE 18th International Symposium on. pp. 25–32. IEEE (2017)
29. Lee, W.S., Grosh, D.L., Tillman, F.A., Lie, C.H.: Fault tree analysis, methods, and applications; a review. IEEE Transactions on Reliability R-34(3), 194–203 (Aug 1985)
30. Lund, M.S., Solhaug, B., Stølen, K.: Model-driven risk analysis: the CORAS approach. Springer Science & Business Media (2010)
31. Mauw, S., Oostdijk, M.: Foundations of attack trees. In: International Conference on Information Security and Cryptology. pp. 186–198. Springer (2005)
32. Meland, P.H., Gjære, E.A.: Representing threats in BPMN 2.0. In: Availability, Reliability and Security (ARES), 2012 Seventh International Conference on. pp. 542–550. IEEE (2012)
33. Meland, P.H., Tøndel, I.A., Jensen, J.: Idea: reusability of threat models–two approaches with an experimental evaluation. In: International Symposium on Engineering Secure Software and Systems. pp. 114–122. Springer (2010)
34. Michel, C.D., Thomas, P.F., Tucci, A.E.: Cyber Risks in the Marine Transportation System. The U.S. Coast Guard Approach
35. Mohr, R.: Evaluating cyber risk in engineering environments: A proposed framework and methodology (2016)
36. Nesheim, D., Rødseth, Ø., Bernsmed, K., Frøystad, C., Meland, P.: Risk model and analysis. Tech. rep., CySIMS (2017)
37. NevilleClarke: Taking-off with BowTie (2013), http://www.nevilleclarke.com/indonesia/articles/topic/52/title/
38. Ni, H., Chen, A., Chen, N.: Some extensions on risk matrix approach. Safety Science 48(10), 1269–1278 (2010)
39. Nielsen, D.S.: The cause/consequence diagram method as a basis for quantitative accident analysis. Tech. rep., Danish Atomic Energy Commission (1971)
40. Phillips, C., Swiler, L.P.: A graph-based system for network-vulnerability analysis. In: Proceedings of the 1998 workshop on New security paradigms. pp. 71–79. ACM (1998)
41. Piètre-Cambacédès, L., Bouissou, M.: Cross-fertilization between safety and security engineering. Reliability Engineering & System Safety 110, 110–126 (2013)
42. Raspotnig, C., Karpati, P., Katta, V.: A Combined Process for Elicitation and Analysis of Safety and Security Requirements, pp. 347–361. Springer Berlin Heidelberg, Berlin, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-31072-0_24
43. Ruijters, E., Stoelinga, M.: Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools. Computer science review 15, 29–62 (2015)
44. Santamarta, R.: A wake-up call for satcom security. Technical White Paper (2014)
45. Schneier, B.: Attack trees. Dr. Dobbs journal 24(12), 21–29 (1999)
46. Sha, L., Gopalakrishnan, S., Liu, X., Wang, Q.: Cyber-physical systems: A new frontier. In: Sensor Networks, Ubiquitous and Trustworthy Computing, 2008. SUTC'08. IEEE International Conference on. pp. 1–9. IEEE (2008)
47. Shostack, A.: Threat modeling: Designing for security (2014)
48. Simon, H.A.: The sciences of the artificial. MIT press (1996)
49. Sindre, G., Opdahl, A.L.: Eliciting security requirements with misuse cases. Requirements engineering 10(1), 34–44 (2005)

50. Sun, M., Mohan, S., Sha, L., Gunter, C.: Addressing safety and security contradictions in cyber-physical systems. In: Proceedings of the 1st Workshop on Future Directions in Cyber-Physical Systems Security (CPSSW09) (2009)
51. Viscusi, W.K., Aldy, J.E.: The value of a statistical life: a critical review of market estimates throughout the world. Journal of risk and uncertainty 27(1), 5–76 (2003)
52. Winther, R., Johnsen, O.A., Gran, B.A.: Security assessments of safety critical systems using hazops. In: International Conference on Computer Safety, Reliability, and Security. pp. 14–24. Springer (2001)
53. Zalewski, J., Drager, S., McKeever, W., Kornecki, A.J.: Towards experimental assessment of security threats in protecting the critical infrastructure. In: ENASE 2012-Proceedings of the 7th International Conference on Evaluation of Novel Approaches to Software Engineering, Wroclaw, Poland (2012)